arXiv:1409.5872v1 [cs.SE] 20 Sep 2014

Incremental Bounded Model Checking for Embedded Software (extended version)*

Peter Schrammel[1], Daniel Kroening[1], Martin Brain[1], Ruben Martins[1], Tino Teige[2], and Tom Bienmüller[2]

[1] University of Oxford
[2] BTC Embedded Systems AG

Abstract. Program analysis is on the brink of mainstream in embedded systems development. Formal verification of behavioural requirements, finding runtime errors and automated test case generation are some of the most common applications of automated verification tools based on Bounded Model Checking. Existing industrial tools for embedded software use an off-the-shelf Bounded Model Checker and apply it iteratively to verify the program with an increasing number of unwindings. This approach unnecessarily wastes time repeating work that has already been done and fails to exploit the power of incremental SAT solving. This paper reports on the extension of the software model checker CBMC to support *incremental Bounded Model Checking* and its successful integration with the industrial embedded software verification tool BTC EMBEDDEDTESTER. We present an extensive evaluation over large industrial embedded programs, which shows that incremental Bounded Model Checking cuts runtimes by *one order of magnitude* in comparison to the standard non-incremental approach, enabling the application of formal verification to large and complex embedded software.

1 Introduction

Recent trend estimation [25] in automotive embedded systems revealed ever growing complexity of computer systems, providing increased safety, efficiency and entertainment satisfaction. Hence, automated design tools are vital for managing this complexity and support the verification processes in order to satisfy the high safety requirements stipulated by safety standards and regulations. Similar to the developments in hardware verification in the 1990s, verification tools for embedded software are becoming indispensable in industrial practice for hunting runtime bugs, checking functional properties and test suite generation [24]. For example, the automotive safety standard ISO 26262 [36] requires the test suite to satisfy modified condition/decision coverage [30] – a goal that is laborious to achieve without support by a model checker that identifies unreachable test goals and suggests test vectors for difficult-to-reach test goals.

* The research leading to these results has received funding from the ARTEMIS Joint Undertaking under grant agreement number 295311 "VeTeSS" and ERC project 280053.

In this paper, we focus on the application of Bounded Model Checking (BMC) to this problem. The technique is highly accurate (no false alarms) and is furthermore able to generate counterexamples that aid debugging and serve as test vectors. The spiralling power of SAT solvers has made this technique scale to reasonably large programs and has enabled industrial application.

In BMC, the property of interest is checked for traces that execute loops up to a given number of times k. Since the value of k that is required to find a bug is not known a-priori, one has to try increasingly larger values of k until a bug is found. The analysis is aborted when memory and runtime limits are exceeded.[3]

Most existing industrial verification tools use an off-the-shelf Bounded Model Checker and, without additional information about the program to be checked, apply it in an iterative fashion:

```
k=0
while true do
  if BMC(program,k) fails then
    return counterexample
  fi
  k++
od
```

This basic procedure offers scope for improvement. In particular, note that the Bounded Model Checker has to redo the work of generating and solving the SAT formula for time frames 0 to k when called to check time frame $k+1$. It is desirable to perform the verification *incrementally* for iteration $k+1$ by building upon the work done for iteration k.

Incremental BMC has been applied successfully to the verification of hardware designs, and has been reported to yield substantial speedups [51,22]. Fortunately, the typical control-loop structure of embedded software resembles the monolithic transition relation of hardware designs, and thus strongly suggests incremental verification of successive loop unwindings. However – to our knowledge – none of the software model checkers for C programs that have competed in the TACAS 2014 Software Verification Competition implement such a technique that ultimately exploits the full power of incremental SAT solving [53,21].

Contributions. The primary contribution of this paper is *experimental*. We quantify the benefit of incremental BMC in the context of the verification of embedded software. To this end,

(1) we survey the techniques that are state of the art in embedded software verification, briefly summarise the underlying theory, and highlight the challenges faced when applying them to industrial code;
(2) we present the first industrial-strength implementation of incremental BMC in a software model checker for ANSI-C programs combining symbolic execution, slicing and incremental SAT solving. Besides loop unwinding, we also elucidate other applications of incremental SAT solving in a state-of-the-art Bounded Model Checker;

[3] One can stop unwinding when the *completeness threshold* [40,39] of the system is reached, but this threshold is often impractically large.

(3) we report on the successful integration of our incremental bounded model checker in the industrial embedded software verification tools BTC EmbeddedTester and EmbeddedValidator where it is used by several hundred industrial users since version 3.4 and 4.3, respectively; and
(4) we give a comprehensive experimental evaluation over a large set of industrial embedded benchmarks that quantify the performance gain due to incremental BMC: our new method outperforms the winner of the TACAS 2014 Software Verification Competition [41] by one order of magnitude.

2 Verification of Model-based Embedded Software

Model-based development is well-established in embedded software development, and particularly popular in the automotive industry. Tools such as Simulink[4] are in widespread use for modelling, code generation and testing.[5]

In this paper, we focus on the verification of C code generated from these models. To this end, we illustrate the characteristics of this verification problem with the help of a well-known case study (Sec. 2.1) and explain the workflow and principal techniques that a state-of-the-art embedded software verification tool uses.

2.1 Case Study: Fault-Tolerant Fuel Control System

The Fault-Tolerant Fuel Control System[6] (FuelSys) for a gasoline engine, originally introduced as a demonstration example for MATLAB Simulink/Stateflow and then adapted for dSPACE TargetLink, is representative of a variety of automotive applications as it combines discrete control logic via Stateflowwith continuous signal flow expressed by Simulink or TargetLink and thus establishes a hybrid discrete-continuous system. More precisely, the control logic of FuelSys is implemented by six automata with two to five states each, while the signal flow is further subdivided into three subsystems with a rich variety of Simulink/TargetLink blocks involving arithmetic, lookup tables, integrators, filters and interpolation (Fig. 1). The system is designed to keep the air-fuel ratio nearly constant depending on the inputs given by a throttle sensor, a speed sensor, an oxygen sensor (EGO) and a pressure sensor (MAP). Moreover it is tolerant to individual sensor faults and is designed to be highly robust, i.e. after detection of a sensor fault the system is dynamically reconfigured.

Properties of interest. The key functional property for FuelSys is how the air-fuel ratio evolves for each of the four sensor-failure scenarios. Simulation-based approaches show that FuelSys is indeed fault-tolerant in each case of a single failure: the air-fuel ratio can be regulated after a few seconds to about

[4] `http://www.mathworks.co.uk/products/simulink/`
[5] The topic of model-based testing methods is discussed in detail in a range of surveys [14,44,46].
[6] `http://www.mathworks.co.uk/help/simulink/examples/modeling-a-fault-tolerant-fuel-control-system.html`

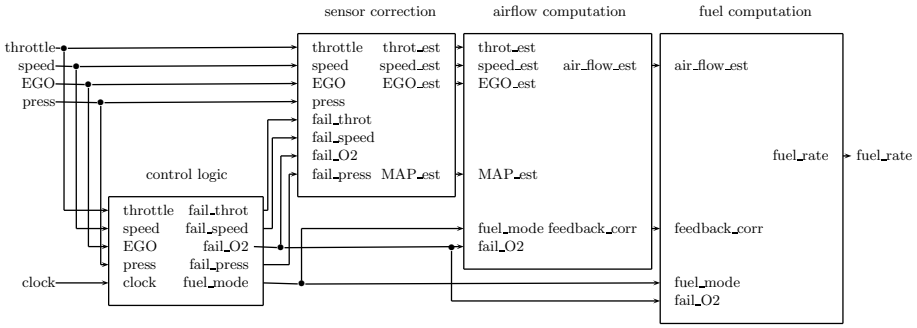

Fig. 1: The SIMULINK Diagram for the Fault-Tolerant Fuel Control System

80 % of the target ratio. In addition to *functional* testing of industrial embedded software, safety standards call for *structural* testing of the production code before release deployment. In Sec. 2.3, we give a brief overview about such standards and the state of affairs of their implementation in practice.

2.2 Structure of Generated Code

Many modelling languages follow the *synchronous programming paradigm* [28], which is well-suited for modelling time-triggered systems, in which tasks (subsystems of the model) execute at given rates. Code generation for such languages produces a typical code structure, which corresponds essentially to a non-preemptive operating system task scheduler. Most code generators provide the scheduler for time-triggered execution or code to interface with popular real-time operating systems. In either case, the functionality corresponds to the following pseudo code:

```
1  void main() {
2    state s; inputs i; outputs o;
3    initialize(s);
4    while(true) { //main loop
5      i = read_inputs();
6      (o,s) = compute_step(i,s);
7      write_outputs(o);
8      wait(); //wait for timer interrupt
9    }
10 }
```

The distinguishing characteristic of such a reactive program is its unbounded main loop, which we will analyse incrementally. All other loops contained within that loop, e.g. to iterate over arrays or interpolate values using look-up tables, have a statically bounded number of iterations and can be fully unwound.

2.3 Analysis with BMC and k-induction

Recent safety standards, e.g. ISO-26262 [36], cover model-based development and testing techniques for early simulation, testing and verification, and recommend

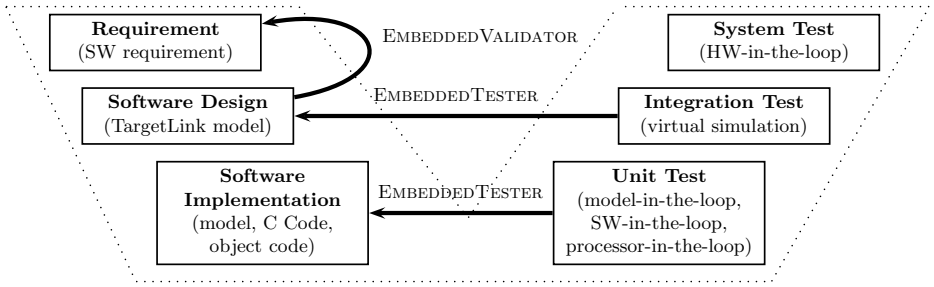

Fig. 2: Embedded software development tool chain in the V model

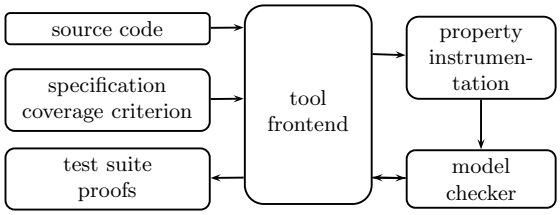

Fig. 3: Typical architecture of a model checker for embedded software

back-to-back testing [42] for showing simulation equivalence between a high-level model and corresponding production code. In the automotive industry, model-based development including automatic code generation is well-established. In particular, MATLAB Simulink by The Mathworks for functional modelling and dSPACE TargetLink[7] for automatic code generation from these models are prominent representatives. Simulink DesignVerifier[8], BTC Embedded-Validator and EmbeddedTester[9], and Reactis[10] complement the software development tool chain as illustrated in Fig. 2 for formal verification of safety requirements against design models and requirement-based testing and back-to-back testing including automatic test vector generation for structural coverage criteria.

Property Instrumentation. Formal verification requires formalisations of high-level requirements, often using observer Büchi automata [12] with a dedicated 'error state' generated from temporal logic descriptions. Test vector generation is done for code-coverage criteria such as branches, statements, conditions and MC/DC [30] of the production C code. For FuelSys, for example, MC/DC instrumentation yields 251 test goals. The properties to be verified or tested have in common that they can be reduced to a reachability problem. In formal

[7] http://www.dspace.com/en/pub/home/products/sw/pcgs/targetli.cfm
[8] http://www.mathworks.com/products/sldesignverifier
[9] http://www.btc-es.de/index.php?lang=2&idcatside=14
[10] http://www.reactive-systems.com

verification of safety properties, we prove that the error state is unreachable, whereas the aim of test vector generation is to obtain a trace that demonstrates reachability of the goal state.

To validate whether the air-fuel ratio in the FuelSys controller is regulated after a few seconds to be within some margin of the target ratio, one has to instrument the reactive program, as sketched above, with an observer implementing the asserted property. For instance, consider the requirement "If some sensor fails for the first time then within 10 seconds the air-fuel ratio will keep in between the range of 80 % to 120 % of the target ratio forever." The code fragment for an observer for this requirement may look as follows:

```
1  // detection of first sensor failure
2  if (sensor_fail == 1 && observe_ratio == 0) {
3      // initialize observer variables
4      observe_ratio = 1;
5      counter = 0;
6      violated = 0;
7  }

8  // if in observation mode...
9  if (observe_ratio == 1) {
10     if (counter >= 10 &&
11         (air_fuel_ratio < 0.8*target_ratio ||
12          air_fuel_ratio > 1.2*target_ratio)) {
13         violated = 1;
14     }
15     counter++;
16 }

17 // safety property
18 assert(violated == 0);
```

In order to verify that the above property actually holds, one has to show that the assertion in the observer code is always satisfied. We use BMC for refutation of the assertion, and k-induction for proving it.

Bounded Model Checking. BMC [3,16] can be used to check the existence of a path $\pi = \langle s_0, s_1, \ldots, s_k \rangle$ of length k between two sets of states described by propositional formulae ϕ to ψ. This check is performed by deciding satisfiability of the following formula using a SAT or SMT solver:

$$\phi(s_0) \wedge \bigwedge_{0 \leq j < k} T(s_j, i_j, s_{j+1}) \wedge \psi(s_k) \tag{1}$$

If the solver returns the answer "satisfiable", it also provides a satisfying assignment to the variables $(s_0, i_0, s_1, i_1, \ldots, s_{k-1}, i_{k-1}, s_k)$. The satisfying assignment represents one possible path $\pi = \langle s_0, s_1, \ldots, s_k \rangle$ from ϕ to ψ and identifies the corresponding input sequence $\langle i_0, \ldots, i_{k-1} \rangle$.

Hence, BMC is useful for refuting safety properties (where ϕ gives the set of initial states and ψ defines the error states) and generating test vectors (where ψ defines the test goal to be covered).

Unbounded Model Checking by k-Induction. BMC can prove reachability, whereas unreachability can be shown using induction. The predicate $\neg\psi$ is an (inductive) invariant, i.e., it holds in all reachable states, if each of the following

two formulae, base case (BC) and induction step (SC), are unsatisfiable:

$$\begin{array}{ll} \text{(BC)} & \phi(s) \wedge \psi(s) \\ \text{(SC)} & \neg\psi(s) \wedge T(s, s') \wedge \psi(s') \end{array} \tag{2}$$

Both formulae can be decided with the help of a SAT or SMT solver.

The property of interest is often not inductive, however, and the check above fails. An option is to strengthen the property, e.g., using auxiliary invariants obtained using an abstract interpreter. Furthermore, the criterion above can be generalised to k-induction [49,22,27,18]:

$$\begin{array}{ll} \text{(BC)} & \phi(s_0) \wedge \bigwedge_{0 \leq j < k} \neg\psi(s_j) \wedge T(s_j, i_j, s_{j+1}) \wedge \psi(s_k) \\ \text{(SC)} & \bigwedge_{0 \leq j \leq k} \neg\psi(s_j) \wedge T(s_j, i_j, s_{j+1}) \wedge \psi(s_{k+1}) \end{array} \tag{3}$$

The base case checks whether $\neg\psi$ holds in the first k steps, whereas the induction step checks if we can conclude from the invariant holding over any k consecutive steps that it holds for the $(k+1)^{st}$ step. If the base step fails, i.e. above formula is satisfiable, we have refuted the property. If it holds and the induction step fails, we do not know whether $\neg\psi$ is invariant. Only if both formulae hold we have proved that $\neg\psi$ is invariant.

Both base step and induction step are essentially instances of BMC: starting from the initial state ϕ for the base case, and starting from *any* state for the induction step. Thus, similar to BMC, k-induction can be applied by using a sequence of increasing values for k.

Challenges. Embedded C code has to meet many conflicting requirements like real-time constraints, low memory footprint and low energy consumption. Code generators offer options to perform certain optimisations towards these goals, often to the detriment of *code size* (and also readability for humans). The observer instrumentation to encode properties and identify the test goals corresponding to code-coverage criteria such as MC/DC produces a non-negligible overhead in the size of the code but introduces little semantic complexity. The size of the formula built further increases whenever internal loops need to be unwound, for example to perform a linear search in lookup tables in the FuelSys example. File sizes of 10 MB and more are common, which poses difficulties to many tools already when parsing the source code and encoding the program into a SAT formula, mostly due to inefficient data structures. Incremental BMC helps reducing formula sizes and peak memory consumption (see Sec. 4.2) by incremental formula generation and solving.

The next generation of automotive electronic control units will be equipped with *floating point* (FP) units [35] to perform complex arithmetic used to implement control laws. FuelSys is an example of this trend. The use of FP arithmetic comes with many hard-to-avoid pitfalls. Many verification tools interpret FP numbers as reals, an approach which gives rise to false positives and is unable to detect certain bugs. Hence, even though bit-precise reasoning about FP arithmetic is expensive, it is indispensable in order to spot intricate bugs. State-of-the-art methods are based on bit-blasting and bitvector refinement [6] (see Sec. 3.3), but new abstraction-based solvers are emerging (e.g. [29][29,5]).

In practice, many loop unwindings may be needed to detect errors and reach certain tests goals (more than 100 for some of our industrial benchmarks, see Sec. 4.2). *Non*-incremental bounded model checking repeats work such as file parsing, loop unwinding, SAT formula encoding and discards information learnt in the SAT solver every time it is called and so gives away an enormous amount of performance. This effect is exacerbating the cost of large unwinding limits that may be needed.

A great challenge is to exploit all the benefits of incrementality in BMC and to significantly enhance performance of its integration with an industrial-strength embedded verification and test-vector generation tool.

3 Incremental BMC

In this section, we explain the technical background of incremental SAT solving and how it is employed in our implementation of incremental BMC.

3.1 Incremental SAT solving

The first ideas for incremental SAT solving date back to the 1990s [34,50,38]. The question is how to solve a sequence of similar SAT problems while reusing effort spent on solving previous instances, i.e., reusing the internal state and learnt information of the solver.

Obviously, incremental SAT solving is easy when the modification to the CNF representation of the problem makes it grow monotonically. This means that if we want to solve a sequence of (increasingly constrained) SAT problems with CNF formulae $\Phi(k)$ for $k \geq 0$ then $\Phi(k)$ must be *growing monotonically* in k, i.e. $\Phi(k+1) = \Phi(k) \wedge \varphi(k)$ for CNF formulae $\varphi(k)$.

Removal of clauses from $\Phi(k)$ is trickier, as some of the clauses learnt during the solving process are no longer implied by the new instance, and need to be removed as well. Earlier work has identified conditions for the reuse of learnt clauses [51,53], but this requires expensive book-keeping, which partially saps the benefit of incrementality.

The most popular approach to incremental solving is to solve SAT problems *under assumptions* [22]: assumptions are modelled as the first decision literals made by the SAT solver. If a learnt clause is derived from an assumption, then it will contain that assumption literal. This light book-keeping enables the SAT solver to maintain its performance when using assumptions. SAT solving under assumptions allows us to emulate the removal of clauses as explained in Sec. 3.2. SMT solvers offer an interface for pushing and popping clauses in a stack-like manner. Pushing adds clauses, popping removes them from the formula. This makes the modification of the formula intuitive to the user, but the efficiency depends on the underlying implementation of the push and pop operations. For example, in [26] it was observed that some SMT solvers (like Z3) are not optimised for incremental usage and hence perform worse incrementally

than non-incrementally. Since CBMC itself implements powerful bitvector decision procedures, we use the SAT solver MINISAT2 [21] as a backend solver.

We will see that incremental BMC requires a *non-monotonic* series of formulae. For SAT solvers, solving under assumptions is the prevalent method, hence we will focus on this technique in the sequel.

3.2 Incremental BMC

Following the construction in [22] for finite state machines, incremental BMC can be formulated as a sequence of SAT problems $\Phi(k)$ that we need to solve:

$$\begin{aligned} \Phi(0) &:= \phi(s_0) \wedge (\Psi(0) \vee \alpha_0) \\ &\quad \text{with assumption } \neg\alpha_0 \\ \Phi(k+1) &:= \Phi(k) \wedge T(s_k, i_k, s_{k+1}) \wedge \alpha_k \wedge \\ &\quad (\Psi(k+1) \vee \alpha_{k+1}) \\ &\quad \text{with assumption } \neg\alpha_{k+1} \end{aligned} \tag{4}$$

where $\Psi(k)$ is the disjunction $\bigvee_{0 \leq j \leq k} \psi(s_j)$ of error states ψ to be proved unreachable up to iteration k. This disjunction means that the verification fails if *at least one* of the error states is reachable. Since the set of ψ_js grows in each iteration, our problem is not monotonic: one has to *remove* $\Psi(k)$ when adding $\Psi(k+1)$ because $\Psi(k)$ subsumes $\Psi(k+1)$, and thus simply conjoining the $\Psi(k)$s would not yield the desired formula.

Here, solving under assumptions comes to rescue. In iteration k, the α_k is assumed to be false, whereas it is assumed true for iterations $k' > k$. This has the effect that in iteration k' the formula $(\Psi(k) \vee \alpha_k)$ becomes trivially satisfied. Hence, it does not contribute to the (un)satisfiability of $\Phi(k')$, which emulates its deletion.[11]

Symbolic execution. For software, however, (4) results in large formulae and would be highly inefficient for the purpose of BMC. In practice, software model checkers use *symbolic execution* in order to exploit, for example, constant propagation and pruning branches when conditionals are infeasible, while generating the SAT formula and thus reducing its size. This means that the formulae describing T and Ψ in (4) are actually dependent on k. Fortunately, this does not affect the correctness of above formula construction and we can replace T by T_k in (4) and ψ by ψ_k in the definition of $\Psi(k)$.

Slicing. Another feature used by state-of-the-art software model checkers is slicing: The purpose of slicing is, again, reducing the size of the SAT formula by removing (or better: not generating) those parts of the formula that have no influence on its satisfiability. There are many techniques how to implement slicing with the desired trade-off between runtime efficiency and its formula pruning effectiveness [52].

[11] For a large number of iterations k, such trivially satisfied subformulas might accumulate as "garbage" in the formula and slow down its resolution. Restarting the solver at appropriate moments is the common solution to this issue.

Slicing is performed relative to $\Psi(k)$. We said that the number of disjuncts ψ_j in Ψ is growing monotonically with k. Hence, we will show that, assuming that our slicing operator is monotonic, we obtain a monotonic formula construction: The transition relation for each time frame T_k is a conjunction $\bigwedge_{\tau \in M} \tau$ of subrelations τ (e.g., formulae corresponding to program instructions). The slicing operator *slice* selects a subset of M. The operator *slice* is monotonic iff $M \subseteq M' \implies slice(M) \subseteq slice(M')$. We can then view the conjunction of transition relations for k time frames $\widehat{T}(k) = \bigwedge_{0 \leq j \leq k} T_j$ as $\bigwedge_{\tau \in M_k} \tau$. Then a slice $\widehat{T}^{sliced}(k)$ of $\widehat{T}(k)$ is $\bigwedge_{\tau \in M'_k} \tau$ where $M'_k \subseteq M_k$. An incremental slice is then defined as the difference between $\widehat{T}^{sliced}(k+1)$ and $\widehat{T}^{sliced}(k)$:

$$T^{sliced}_{k+1} = \bigwedge_{\tau \in M'_{k+1} \setminus M'_k} \tau. \tag{5}$$

Monotonicity of formula construction follows from $M'_{k+1} \subseteq M_{k+1}$ and the assumed monotonicity $M'_k \subseteq M'_{k+1}$ of the slicing operator. We can thus replace T by T^{sliced}_k in (4). Mind that T^{sliced}_k contains also subrelations τ for time steps $k' < k$.

3.3 Incremental Refinements

Incremental SAT solving is also used for incremental refinements of the transition relation T for bitvectors and arrays, for example.

Bitvector Refinement. The purpose of bitvector refinement [10,2,31,11,9,19] is to reduce the size of formulae encoding bitvector operations. This is especially important for arithmetic operations that generate huge SAT formulae, e.g. multiplication, division and remainder operations, both for integer and floating-point variables [6]. Bitvector refinement is based on successive under- and overapproximations. For instance, underapproximations can be obtained by fixing a certain number of bits, whereas overapproximation make a certain number of bits unconstrained. If an underapproximation is satisfiable (SAT) or an overapproximation is unsatisfiable (UNSAT) we know that the non-approximated formula is SAT or UNSAT respectively. Otherwise, the number of fixed respectively unconstrained bits is reduced until the non-approximated formula itself is checked.

Arrays. To handle programs with arrays, Ackermann expansion is necessary to ensure the functional consistency property of arrays: $\forall i, j : i = j \implies A[i] = A[j]$. However, adding a quadratic number of constraints (in the size of the array A) is extremely costly. Experience has shown that only a small number of these constraints is actually used [47].

Hence, more efficient just trying to solve the SAT formula without these constraints, which is an over-approximation. Hence, if we get an UNSAT result (a), we know that the solution with the Ackermann constraints would be UNSAT too. In case of a SAT result (b), we check the consistency of the obtained model: if it turns out not to violate consistency, then we know that we have found a real bug. Otherwise (c), we add the violated Ackermann constraint to the formula.

The formula construction is trivially monotonic and we can use incremental SAT solving. We repeat the procedure until we hit case (a) or (b), which is guaranteed to happen.

Some SMT solvers, like BOOLECTOR implement a similar procedure to decide the SMT-LIB array theory [7,8].

Formula construction. Applying above refinements inside an incremental Bounded Model Checker requires using several incremental formula encodings for (in general, non-monotonic) refinements simultaneously. These refinements are global over all unwindings, so that in iteration k we have to further refine transition relations $T_{k'}$ from earlier iterations $k' < k$. We can formalise the incremental formula construction as follows: For iteration $k \geq 0$ of incremental BMC and the ℓ^{th} refinement:

$$\begin{aligned}
\Phi(0,0) &:= \phi(s_0) \wedge (\Psi_0(s_0) \vee \alpha_0) \\
&\quad \text{with assumption } \neg\alpha_0 \\
\Phi(k+1,\ell) &:= \Phi(k,\ell) \wedge (\Psi_{k+1}(s_{k+1}) \vee \alpha_{k+1}) \wedge \alpha_k \wedge \\
&\quad \big(T'_{k+1,\ell}(s_k, i_k, s_{k+1}) \vee \beta_\ell\big) \\
&\quad \text{with assumptions } \neg\alpha_{k+1} \text{ and } \neg\beta_\ell
\end{aligned}$$

$$\begin{aligned}
\Phi(k,\ell+1) &:= \Phi(k,\ell) \wedge \\
&\quad \big(T'_{k-1,\ell+1}(s_{k-1}, i_{k-1}, s_k) \vee \beta_{\ell+1}\big) \wedge \beta_\ell \\
&\quad \text{with assumptions } \neg\alpha_k \text{ and } \neg\beta_{\ell+1} \\
&\quad \text{for } k \geq 1
\end{aligned}$$

ℓ is incremented in each iteration of the refinement loop until convergence, whereas k is incremented when considering the next time frame.

4 Experimental Evaluation

We present the results of our experimental evaluation of incremental BMC and incremental k-induction on industrial programs from mainly automotive origin. The experiments for this study were performed on a 3.5 GHz Intel Xeon machine with 8 cores and 32 GB of physical memory running Windows 7 with a time limit of 3,600 seconds.

Our study of incremental BMC is targeted at embedded software since it takes advantage of its specific properties (one big unbounded loop, whereas other loops are bounded). However, incremental BMC can also be applied to programs where loops and control structures are more irregular; to this end, we report on a preliminary study of incremental BMC on programs with multiple loops. These latter experiments were performed on a 3 GHz Intel Xeon with 8 cores and 50 GB of physical memory running Fedora 20 with a time limit of 3,600 seconds.

4.1 Implementation

Cbmc. We have implemented our extension[12] for incremental BMC in the Bounded Model Checker for ANSI-C programs CBMC [17] using the SAT solver MINISAT2 [21] as a backend solver.

CBMC is called in incremental mode using the command line `cbmc file.c --no-unwinding-assertions --incremental`.[13] The following options can be added to enable specific features of CBMC:

- `--no-sat-preprocessor`: turns off SAT formula preprocessing, i.e. the MINISAT2 simplifier is not used.
- `--slice-formula`: slices the SAT formula.
- `--refine`: enables bitvector refinement.
- `--unwind-max` k: limits the loop unwindings of the loop to be checked incrementally to k unwindings. Without this option, CBMC will not terminate for unsatisfiable instances, i.e. bug-free programs with unbounded loops.

More information regarding the usage of incremental CBMC can be found on the CPROVER wiki page.[14]

Integration with an industrial-strength embedded verification tool. In the integration of CBMC with BTC EMBEDDEDTESTER and EMBEDDEDVALIDATOR, a master routine selects the next verification/test goal to be analysed starting from instrumented C code. After some preprocessing like source-level slicing and internal-loop unwinding the resulting reachability task is given to CBMC. If CBMC is able to solve the problem within the user-defined time limit, the result, i.e. information of bounded or unbounded unreachability, or a test vector or counterexample in case of reachability, is reported back to the master process. Otherwise, i.e. in case of a timeout, CBMC is killed but information about the solved unwindings of the reactive main loop is given back, which frequently is a useful result for the user.

To prove unreachability of verification/test goals (properties), split-case k-induction is performed (see Section 2.3). For this purpose BTC EMBEDDEDTESTER generates two source files, one containing the base case, which is a normal BMC problem with the property given as assertion (cf. Equ. (3) (BC)); the file for the step case havocs variables modified in the loop and the invariant property is assumed at the beginning of the loop and asserted at the end of the loop (cf. Equ. (3) (SC)). To check the step case, we require a reversed termination behaviour of CBMC (option `--stop-when-unsat`), i.e. it continues unwinding as long as the problem is SAT and stops as soon as it is UNSAT.

Implementation of Incremental BMC for General Sequential Programs.

[12] Source code available from `http://www.cprover.org/svn/cbmc/branches/peter-incremental-unwinding`

[13] The option `--no-unwinding-assertions` is required to prevent CBMC from inserting an assertion that makes the verification fail when a loop has not been fully unwound, which is useless in our application with unbounded loops.

[14] `http://www.cprover.org/wiki/doku.php?id=how_to_use_incremental_unwinding`

			operators			input variables			state variables			observer	
		LOC	cond	mul	div/rem	bool	int	float	bool	int	float	bool	unwindings
SAT	max	31222	17103	669	75	688	477	189	3876	750	107	22	106
	average	7572	4306	188	9	103	79	19	583	136	15	9	22
UNSAT	max	23014	49530	567	37467	212	282	188	708	663	32	22	10
	average	4854	6014	160	1257	30	51	9	163	73	3	7	10

Table 1: Benchmark characteristics from industrial programs

Incremental CBMC can also be used for programs with multiple loops. For these programs, CBMC incrementally unwinds loops one at each time. For a given loop, CBMC will unwind the loop until it is fully unwound or until a maximum depth k is reached (given by option `--unwind-max` k). After a loop has been unwound, CBMC continues to the next loop. This procedure is repeated until all loops have been unwound or a bug has been found. Recursive function calls are treated similarly.

4.2 Incremental BMC for Embedded Software

We report results on industrial programs for the integration of CBMC with BTC EMBEDDEDTESTER and EMBEDDEDVALIDATOR. For these experiments, we used 60 benchmarks that originated mainly from automotive applications.[15] Half of the benchmarks are bug-free (UNSAT instances), half contain a bug (SAT instances). This benchmark suite is an indicator for performance of model checking tools in practice as it covers a representative spectrum of embedded software. A summary of the benchmark characteristics is listed in Table 1; for a full list we refer to Table 2 in the Appendix.

Besides the number of lines of code, we give the number of conditional operators, multiplications and divisions or remainder operations, which are a good indicator for the difficulty of the benchmark, because they generate large formulae — recall that for each "/" occurring in the program, CBMC has to generate a divider circuit. The surprisingly high number of conditional operators in most of the benchmarks is due to the preprocessing of conditional assignments by BTC EMBEDDEDTESTER and hints at the amount of branching in these benchmarks. Moreover, we list the number of input and state variables, and the variables introduced by the observer instrumentation.

These benchmarks have the property of having only one unbounded loop. For these benchmarks, CBMC is called in incremental mode by using the option `--incremental-check c:main.0` where `c::main.0` is the loop identifier of the unbounded loop to be unwound and checked incrementally. The loop identifiers can be obtained using the option `--show-loops`.

Runtimes. We compared the incremental (i) with the non-incremental (ni) approach and evaluated the impact of slicing (s), SAT preprocessing (p) and bitvec-

[15] Unfortunately, the source code of the benchmarks cannot be made public due to non-disclosure agreements.

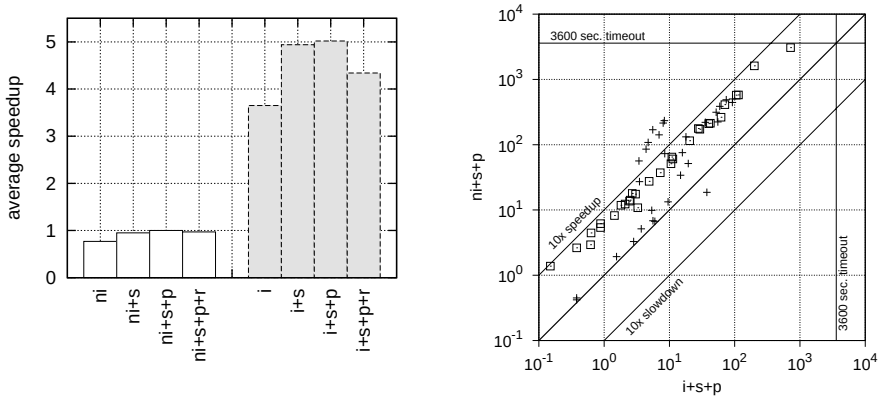

(a) Effect of slicing, SAT formula preprocessing and bitvector refinement

(b) Comparison between ni+s+p and i+s+p (+ SAT instances; □ UNSAT instances)

Fig. 4: Incremental vs. non-incremental BMC

tor refinement (r).[16] The incremental and non-incremental approaches were compared by activating none of the three techniques, with slicing only (+s), with slicing and preprocessing (+s+p), and with all three options activated (+s+p+r).

The maximum number of loop unwindings was fixed to 10 for the UNSAT instances in order to balance a significant exploration depth with reasonable analysis runtimes. For SAT instances, a maximum number of loop unwindings was not fixed since the incremental and non-incremental approaches are bound to terminate when the unwinding depth reaches the depth of the bug. The number of unwindings are listed in the last column in Table 1 (resp. Table 2 in the Appendix).

Fig. 4 shows the comparison between the incremental and non-incremental approaches and the impact of each tool option on their performance.

Fig. 4a shows the average geometric mean [23] speedup of instances that were solved by all approaches. We consider as baseline the (ni+s+p) approach since it was the best non-incremental approach. Each bar shows the average geometric mean speedup of each approach when compared to (ni+s+p). For example, (ni) has a speedup of 0.77, i.e. (ni) is on average 0.77× slower than (ni+s+p). On the other hand, all incremental versions are much faster than the non-incremental versions. For example, (i) is on average over 3.5× faster than (ni+s+p) and (i+s+p) is on average over 5× faster than (ni+s+p). We observe the following effects of the tool options: (i) slicing shows significant benefits overall (also on peak memory consumption); (ii) not using formula preprocessing is a bad idea in general; and (iii) bitvector refinement shows benefits for UNSAT instances, but produces overhead for SAT instances which deteriorates the overall performance of the tool (see Fig. 7 in the Appendix for more details). Even though the tool

[16] Array refinement is not used because the benchmarks do not contain arrays.

options have some positive effects, they are rather minor in comparison to the performance gains from using an incremental approach.

Since the best incremental and non-incremental approaches were obtained with the configuration (+s+p), we will use this configuration for both approaches on the results described in the remainder of the paper.

Fig. 4b shows a scatter plot with runtimes of the best non-incremental (ni+s+p) and incremental (i+s+p) approaches. Each point in the plot corresponds to an instance, where the x-axis corresponds to the runtime required by the incremental approach and the y-axis corresponds to the runtime required by the non-incremental approach. If an instance is above the diagonal, then it means that the incremental approach is faster than the non-incremental approach, otherwise it means that the non-incremental approach is faster. SAT instances are plotted as crosses, whereas UNSAT instances are plotted as squares. Incremental BMC significantly outperforms non-incremental BMC. For SAT instances, the advantage of incremental BMC is negligible for the easy instances, whereas speedups are around a factor of 10 for the medium and hard instances. For UNSAT instances, speedups are also significant and most instances have a speedup of more than a factor of 5.

Solving vs. overall runtime. Since the non-incremental approach has to reparse source files and preprocess them at each iteration, one might argue that removing this overhead is the main reason for the speedup observed. However, the overhead for parsing files, symbolic execution and slicing when compared to generating and solving SAT formula is similar for the incremental and non-incremental approach. The incremental approach spends 27% of its time solving the SAT formula (582 out of 2,151 seconds), whereas the non-incremental approach spends 28% of its time (3,317 out of 11,811 seconds). Unsurprisingly, solving the instance for the largest k in the non-incremental approach takes a considerable amount of time (around 24%), when compared to the total time for solving the SAT formulae for iterations 1...k (784 out of 3,317 seconds).

An explanation for these speedups might be the size of the queries issued in both approaches. The average number of clauses per solver call is halved from 1,367k clauses for the non-incremental approach to 709k clauses for the incremental approach. Similarly, the average number of variables is less than a third in the incremental approach when compared to the non-incremental approach, being of 217k and 746k respectively.

Peak memory consumption. Smaller query sizes also have an effect on peak memory consumption which is reduced by 30% for UNSAT benchmarks; for SAT benchmarks, however, we observed a 10% increase.

4.3 Incremental BMC for Code coverage on FuelSys

As reported in the previous section, enabling Cbmc to work incrementally led to tremendous performance gains on the benchmark suite consisting of selected single input files. In order to assess whether these improvements have practical impact in the *integration* of Cbmc with an industrial-strength test-vector generation tool, we compared the performance of BTC EmbeddedTester with

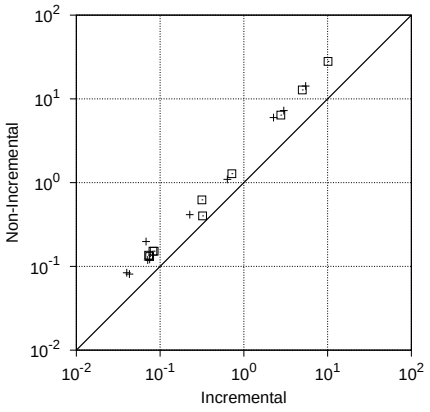

Fig. 5: Incremental k-induction (+ base case; □ step case)

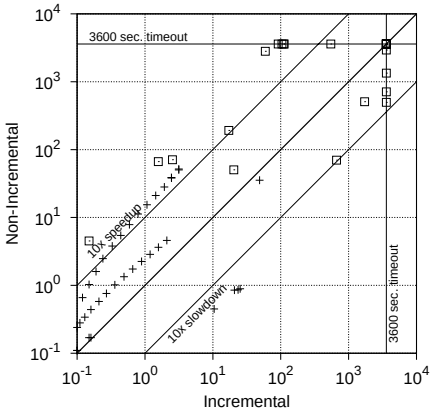

Fig. 6: Incremental BMC for SystemC (+ SAT instances; □ UNSAT instances)

the incremental feature of CBMC being disabled and enabled. The time limit per subtask was 10 minutes and the unwinding depth for all internal loops was 50. For unwinding depth 10 of the main loop, the incremental feature improves the overall runtime from 152.3 to 70.4 minutes, i.e. more than 2× faster, and for unwinding depth 50 from 377.4 to 108.5 minutes, i.e. more than 3× faster. In the latter case, the rate of solved subproblems for MC/DC (i.e. not run into timeout) could be increased from 98.4 % to 99.2 %.

4.4 Incremental k-Induction for Embedded Software

To compare the performance of incremental and non-incremental approaches for k-induction, we considered the subset of UNSAT benchmarks for which k-induction required more than 1 iteration. Note that when k-induction requires only 1 iteration, the performance of incremental and non-incremental approaches is similar. This subset of benchmarks corresponds to 10 UNSAT benchmarks (see Table 2 in the Appendix for more details).

Fig. 5 shows a scatter plot with the runtimes of incremental and non-incremental k-induction using the tool options (+s+p). Instances that correspond to the base case are plotted as crosses, whereas instances that correspond to the step case are plotted as squares. The runtimes for both incremental and non-incremental checking are relatively small. These are due to the small number of iterations required by k-induction to prove the unreachability of the properties present on these benchmarks (between 2 and 4 iterations with an average of 2.4 per instance).

Incremental checking is always faster than non-incremental checking. When considering the average geometric mean speedup, incremental checking is around 2× faster than non-incremental checking, on both base and step cases.

4.5 Incremental BMC for Programs with Multiple Loops

Incremental BMC is not restricted to programs with a single unbounded loop and may also be applied to programs with multiple unbounded loops. To evaluate the performance of incremental BMC on this kind of programs, we compared the performance of incremental and non-incremental approaches on the 62 benchmarks from the SystemC category of the Software Verification Competition benchmark set,[17] because these benchmarks, which were derived from SystemC models [15], contain many loops. 25 benchmarks are bug-free (UNSAT instances) and 37 contain a bug (SAT instances). These benchmarks have between 2 and 19 loops with an average of 10.3 loops per instance. For SAT instances, the depth of the bug ranges from 1 to 5 with an average depth of 2.5. For more details regarding these benchmarks see Table 3 in the Appendix.

We have fixed the maximum number of loop unwindings to 10 for both, SAT and UNSAT instances. Note that this unwind depth is larger than the depth of the bugs for the SAT instances. Formula slicing is not yet fully supported in incremental CBMC for programs with multiple loops. Therefore, the incremental approach was run with the tool options (+p), whereas the non-incremental approach was run with the tool options (+s+p).

Fig. 6 shows a scatter plot with the run times of the incremental and non-incremental approaches. For the majority of the instances, the incremental approach outperforms the non-incremental approach and for many SAT and UNSAT instances the speedup is larger than a factor of 10. However, there are a few instances for which the non-incremental approach performs better. The non-incremental approach unwinds all loops until a fixed unwind depth, whereas the incremental approach fully unwinds one loop before continuing to the next loop. For some instances, fully unwinding each loop may result in the generation of larger formulae, particularly for SAT instances. Not using slicing for the incremental approach may also result in larger formulae. The increase in formula size may explain the observed slowdown for some instances.

Overall, when considering instances solved by both approaches, the incremental approach is faster than the non-incremental approach and the average geometric speedup is larger than a factor of 3.

5 Related Work

Most related is recent work on a prototype tool NBIS [26] implementing incremental BMC using SMT solvers. They show the advantages of incremental software BMC. However, they do not consider industrial embedded software and have evaluated their tool only on small benchmarks that are very easy for both, incremental and non-incremental, approaches (runtimes <1s).[18]

[17] Available at `https://svn.sosy-lab.org/software/sv-benchmarks/trunk/c/systemc/`

[18] Unfortunately, a working version of the tool was not available at time of submission.

Bit-precise formal verification techniques are indispensable for embedded system models and implementations, that have low-level, i.e. C language, semantics like discrete-time Simulink models. The importance of this topic has recently attracted attention as shown by publications on verification using SMT Solving [32,43], test case generation [45], symbolic analysis for improving simulation coverage [1], and directed random testing [48]. Yet, all these works have not exploited incremental BMC.

The test vector generation tool FShell [33] uses incremental SAT solving to check the reachability of a set of test goals. However, it assumes a fixed unwinding of the loops. There is no reason why incremental BMC should not boost its performance when increasing loop unwindings need to be considered. Test vector generation tools like Klee [13] use incremental SAT solving to extend the paths to be explored. However, they consider only single paths at a time, whereas BMC explores all paths simultaneously.

Incremental SAT solving has important applications in other verification techniques like the IC3 algorithm [4,20] and incremental BMC is standard for hardware verification [37,54]. We show that the speedups of incremental SAT solving reported in [22] regarding k-induction on small HW circuits carry over to industrial embedded software.

6 Conclusions

We claim that incremental BMC is an indispensable technique for embedded software verification that should be considered state-of-the-art in such tools. To underpin this claim, we report on the successful integration of our incremental extension of Cbmc into an industrial embedded software verification tool. Our experiments demonstrate one-order-of-magnitude speedups from incremental approaches on industrial embedded software benchmarks for BMC and k-induction. These performance gains result in faster property verification and higher test coverage, and thus, a productivity increase in embedded software verification.

Moreover, we reported on significant speedups for programs with multiple loops that show the applicability of incremental BMC beyond embedded software. We expect that incremental BMC can be further improved for programs with multiple loops by simultaneously unwinding all loops incrementally instead of fully unwinding one loop at each time. Incremental k-induction for programs with multiple loops will also benefit from such an improvement.

References

1. Alur, R., Kanade, A., Ramesh, S., Shashidhar, K.C.: Symbolic analysis for improving simulation coverage of Simulink/Stateflow models. In: EMSOFT. pp. 89–98 (2008)
2. Biere, A.: PicoSAT essentials. JSAT 4(2-4), 75–97 (2008)

3. Biere, A., Cimatti, A., Clarke, E.M., Zhu, Y.: Symbolic model checking without BDDs. In: TACAS. LNCS, vol. 1579, pp. 193–207 (1999)
4. Bradley, A.R.: IC3 and beyond: Incremental, Inductive Verification. In: CAV. LNCS, vol. 7358, p. 4 (2012)
5. Brain, M., D'Silva, V., Griggio, A., Haller, L., Kroening, D.: Interpolation-based verification of floating-point programs with abstract CDCL. In: SAS. LNCS, vol. 7935 (2013)
6. Brillout, A., Kroening, D., Wahl, T.: Mixed abstractions for floating-point arithmetic. In: FMCAD. pp. 69–76 (2009)
7. Brummayer, R., Biere, A.: Boolector: An efficient SMT solver for bit-vectors and arrays. In: TACAS. LNCS, vol. 5505, pp. 174–177 (2009)
8. Brummayer, R., Biere, A.: Lemmas on Demand for the Extensional Theory of Arrays. JSAT 6(1-3), 165–201 (2009)
9. Brummayer, R., Biere, A.: Effective bit-width and under-approximation. In: EUROCAST. LNCS, vol. 5717, pp. 304–311 (2009)
10. Bryant, R.E., Kroening, D., Ouaknine, J., Seshia, S.A., Strichman, O., Brady, B.A.: Deciding Bit-Vector Arithmetic with Abstraction. In: TACAS. LNCS, vol. 4424, pp. 358–372 (2007)
11. Bryant, R.E., Kroening, D., Ouaknine, J., Seshia, S.A., Strichman, O., Brady, B.A.: An abstraction-based decision procedure for bit-vector arithmetic. Journal on Software Tools for Technology Transfer 11(2), 95–104 (2009)
12. Buechi, J.R.: On a Decision Method in Restricted Second-Order Arithmetic. In: International Congress on Logic, Methodology, and Philosophy of Science. pp. 1–11. Stanford University Press (1962)
13. Cadar, C., Dunbar, D., Engler, D.: KLEE: Unassisted and Automatic Generation of High-Coverage Tests for Complex Systems Programs. In: OSDI. pp. 209–224 (2008)
14. Chakraborty, S., Ramesh, S., Teich, J.: Model-based analysis, synthesis and testing of automotive hardware/software architectures. In: EMSOFT. pp. 299–300 (2010)
15. Cimatti, A., Micheli, A., Narasamdya, I., Roveri, M.: Verifying SystemC: A software model checking approach. In: FMCAD. pp. 51–59 (2010)
16. Clarke, E., Biere, A., Raimi, R., Zhu, Y.: Bounded model checking using satisfiability solving. Formal Methods in System Design 19(1), 7–34 (2001)
17. Clarke, E., Kroening, D., Lerda, F.: A tool for checking ANSI-C programs. In: TACAS. LNCS, vol. 2988, pp. 168–176 (2004)
18. Donaldson, A., Haller, L., Kroening, D., Rümmer, P.: Software Verification Using k-Induction. In: SAS. LNCS, vol. 6887, pp. 351–368 (2011)
19. Eén, N., Mishchenko, A., Amla, N.: A single-instance incremental SAT formulation of proof- and counterexample-based abstraction. In: FMCAD. pp. 181–188 (2010)
20. Eén, N., Mishchenko, A., Brayton, R.K.: Efficient implementation of property directed reachability. In: FMCAD. pp. 125–134 (2011)
21. Eén, N., Sörensson, N.: An Extensible SAT-solver. In: SAT. LNCS, vol. 2919, pp. 502–518 (2003)
22. Eén, N., Sörensson, N.: Temporal induction by incremental SAT solving. ENTCS 89:4, 543–560 (2003)
23. Fleming, P., Wallace, J.: How Not To Lie With Statistics: The Correct Way To Summarize Benchmark Results. CACM 29(3), 218–221 (1986)
24. Fraser, G., Wotawa, F., Ammann, P.: Testing with model checkers: a survey. STVR 19(3), 215–261 (2009)
25. Gunnarsson, D., Kuntz, S., Farrall, G., Iwai, A., Ernst, R.: Trends in automotive embedded systems. In: CODES+ISSS. pp. 9–10 (2012)

26. Günther, H., Weissenbacher, G.: Incremental bounded software model checking. In: SPIN. pp. 40–47 (2014)
27. Hagen, G., Tinelli, C.: Scaling up the formal verification of Lustre programs with SMT-based techniques. In: FMCAD. pp. 1–9 (2008)
28. Halbwachs, N.: Synchronous programming of reactive systems. Kluwer (1993)
29. Haller, L., Griggio, A., Brain, M., Kroening, D.: Deciding floating-point logic with systematic abstraction. In: FMCAD (2012)
30. Hayhurst, K.J., Veerhusen, D.S., Chilenski, J.J., Rierson, L.K.: A practical tutorial on modified condition/decision coverage. Tech. rep., NASA (May 2001)
31. He, N., Hsiao, M.S.: A new testability guided abstraction to solving bit-vector formula. In: International Workshop on Bit-Precise Reasoning (2008)
32. Herber, P., Reicherdt, R., Bittner, P.: Bit-precise formal verification of discrete-time MATLAB/Simulink models using SMT solving. In: EMSOFT. pp. 1–10 (2013)
33. Holzer, A., Schallhart, C., Tautschnig, M., Veith, H.: Query-driven program testing. In: VMCAI. LNCS, vol. 5403, pp. 151–166 (2009)
34. Hooker, J.N.: Solving the incremental satisfiability problem. JLP 15(1&2), 177–186 (1993)
35. Infineon Technologies AG: 32-bit TriCore microcontroller
36. ISO 26262: Road vehicles – Functional safety (2011)
37. Jin, H., Somenzi, F.: An incremental algorithm to check satisfiability for bounded model checking. ENTCS 119:2, 51–65 (2005)
38. Kim, J., Whittemore, J., Sakallah, K.A., Silva, J.P.M.: On applying incremental satisfiability to delay fault testing. In: DATE. pp. 380–384 (2000)
39. Kroening, D., Ouaknine, J., Strichman, O., Wahl, T., Worrell, J.: Linear completeness thresholds for bounded model checking. In: CAV. LNCS, vol. 6806, pp. 557–572 (2011)
40. Kroening, D., Strichman, O.: Efficient computation of recurrence diameters. In: VMCAI. LNCS, vol. 2575, pp. 298–309 (2003)
41. Kroening, D., Tautschnig, M.: CBMC – C bounded model checker – (competition contribution). In: TACAS. LNCS, vol. 8413, pp. 389–391. Springer (2014)
42. Liggesmeyer, P.: Software-Qualität - Testen, Analysieren und Verifizieren von Software (2. Aufl.). Spektrum Akademischer Verlag (2009)
43. Manamcheri, K., Mitra, S., Bak, S., Caccamo, M.: A step towards verification and synthesis from Simulink/Stateflow models. In: HSCC. pp. 317–318 (2011)
44. Neto, A.C.D., Travassos, G.H.: A picture from the model-based testing area: Concepts, techniques, and challenges. Advances in Computers 80, 45–120 (2010)
45. Peranandam, P., Raviram, S., Satpathy, M., Yeolekar, A., Gadkari, A.A., Ramesh, S.: An integrated test generation tool for enhanced coverage of Simulink/Stateflow models. In: DATE. pp. 308–311 (2012)
46. Petrenko, A., da Silva Simão, A., Maldonado, J.C.: Model-based testing of software and systems: recent advances and challenges. STTT 14(4), 383–386 (2012)
47. Pnueli, A., Strichman, O.: Reduced functional consistency of uninterpreted functions. ENTCS 144(2), 53–65 (2006)
48. Satpathy, M., Yeolekar, A., Ramesh, S.: Randomized directed testing (REDIRECT) for Simulink/Stateflow models. In: EMSOFT. pp. 217–226 (2008)
49. Sheeran, M., Singh, S., Stålmarck, G.: Checking safety properties using induction and a SAT-solver. In: FMCAD. LNCS, vol. 1954, pp. 108–125 (2000)
50. Silva, J.M., Sakallah, K.A.: Robust search algorithms for test pattern generation. In: FTCS. pp. 152–161 (1997)
51. Strichman, O.: Pruning techniques for the SAT-based bounded model checking problem. In: CHARME. LNCS, vol. 2144, pp. 58–70 (2001)

52. Tip, F.: A survey of program slicing techniques. Tech. rep., CWI-Amsterdam (1994)
53. Whittemore, J., Kim, J., Sakallah, K.A.: SATIRE: A new incremental satisfiability engine. In: DAC. pp. 542–545 (2001)
54. Wieringa, S.: On incremental satisfiability and bounded model checking. In: Design & Impl. of Formal Tools & Sys. pp. 46–54 (2011)

A Industrial Benchmark Characteristics

		operators			input variables			state variables			observer	
name	LOC	cond	mul	div/rem	bool	int	float	bool	int	float	bool	unwindings
automotive_sat_01	3762	2032	82	1	14	282	0	229	50	0	3	12
automotive_sat_02	1854	189	79	1	78	4	0	165	7	0	3	15
automotive_sat_03	15277	17103	669	75	230	244	0	868	275	0	1	9
automotive_sat_04	13853	16908	601	59	208	219	0	741	266	0	1	12
automotive_sat_05	469	193	90	11	1	0	0	17	3	0	3	21
automotive_sat_06	10702	5117	646	1	7	54	19	28	60	22	16	5
automotive_sat_07	10970	5068	646	1	7	54	19	27	62	22	15	4
automotive_sat_08	3656	2657	79	1	14	61	26	20	68	30	16	2
automotive_sat_09	253	34	79	1	0	3	0	23	4	0	3	103
automotive_sat_10	604	117	79	1	23	7	0	81	10	0	3	40
automotive_sat_11	592	115	79	1	23	7	0	79	10	0	3	48
automotive_sat_12	1978	2201	79	1	0	0	0	4	172	0	3	53
automotive_sat_13	1980	2198	79	1	0	0	0	4	172	0	3	55
automotive_sat_14	1222	216	79	1	0	26	0	94	67	0	3	56
automotive_sat_15	5020	3172	79	1	18	4	0	115	22	0	3	17
automotive_sat_16	2578	4572	89	4	1	20	105	3	22	107	17	2
automotive_sat_17	2580	4592	89	4	1	20	105	2	22	107	18	1
automotive_sat_18	2740	4718	89	4	1	20	105	2	24	107	16	2
automotive_sat_19	27456	3579	177	7	546	95	0	3426	438	0	1	12
automotive_sat_20	27456	3579	177	7	546	95	0	3426	438	0	1	16
automotive_sat_21	31222	3705	178	7	688	477	0	3876	750	0	1	12
automotive_sat_22	30834	3620	177	7	652	476	0	3837	744	0	1	14
automotive_sat_23	1270	508	102	5	6	66	0	79	124	9	16	1
automotive_sat_24	1272	501	102	5	6	66	0	78	124	9	17	3
automotive_sat_25	1282	506	102	5	6	67	0	79	128	9	15	1
automotive_sat_26	321	28	79	1	6	2	0	36	2	0	3	106
avionics_sat	2214	1413	79	2	30	16	0	189	52	0	1	20
fuelsys_sat_01	9402	16603	311	6	0	0	4	31	5	8	22	1
fuelsys_sat_02	9404	16757	311	6	0	0	4	31	5	8	22	1
fuelsys_sat_03	5746	8521	224	3	0	0	4	30	5	7	19	1
automotive_unsat_01*	3761	2032	82	1	14	282	0	229	50	0	3	10
automotive_unsat_02	3762	2032	82	1	14	282	0	229	50	0	3	10
automotive_unsat_03	1579	889	79	1	0	38	0	75	4	0	3	10
automotive_unsat_04	1853	189	79	1	78	4	0	165	7	0	3	10
automotive_unsat_05	503	321	106	19	1	0	0	21	3	0	3	10
automotive_unsat_06	13259	16672	545	59	188	207	0	708	232	0	1	10
automotive_unsat_07	464	193	90	11	1	0	0	17	3	0	3	10
automotive_unsat_08	23014	49530	536	37467	92	220	0	697	304	0	1	10
automotive_unsat_09	4768	3334	79	1	0	26	0	215	663	0	3	10
automotive_unsat_10	1035	160	79	1	30	4	0	115	29	0	1	10
automotive_unsat_11	12142	5859	567	0	7	54	19	27	60	22	17	10
automotive_unsat_12	12518	6242	567	0	7	54	19	27	62	22	15	10
automotive_unsat_13*	4726	3091	42	0	14	61	26	30	71	32	16	10
automotive_unsat_14*	591	115	79	1	23	7	0	79	10	0	3	10
automotive_unsat_15*	1977	2198	79	1	0	0	0	4	172	0	3	10
automotive_unsat_16	2339	559	82	9	22	56	0	170	79	0	3	10
automotive_unsat_17*	1399	258	79	1	0	29	0	106	73	0	3	10
automotive_unsat_18*	5021	3172	79	1	18	4	0	115	22	0	3	10
automotive_unsat_19*	7979	12127	119	15	0	0	0	5	16	0	3	10
automotive_unsat_20*	6217	686	88	2	212	87	0	697	60	0	1	10
automotive_unsat_21*	5230	1043	81	2	99	24	0	511	112	0	1	10
automotive_unsat_22	190	97	90	11	0	0	0	4	31	0	1	10
automotive_unsat_23	659	93	79	1	9	1	0	75	10	0	3	10
automotive_unsat_24	3554	787	81	52	16	79	0	226	45	0	3	10
automotive_unsat_25	1575	184	79	1	38	0	0	199	15	0	3	10
avionics_unsat	2329	1413	79	2	30	16	0	188	52	0	1	10
fuelsys_unsat_01*	5146	17271	214	5	0	0	3	11	0	5	21	10
fuelsys_unsat_02	7806	19764	215	6	0	0	4	31	5	8	22	10
fuelsys_unsat_03	7804	19764	215	5	0	0	4	31	5	8	22	10
fuelsys_unsat_04	3340	11671	205	3	0	0	3	15	0	3	18	10

Table 2: Embedded software benchmark characteristics (name of the benchmark and application domain, lines of code, number of operators (cond(`a?b:c`), mul(`*`), div/rem(`/`,`%`)), number of boolean/integer/floating point input and state variables, number of boolean variables introduced by the observer instrumentation, number of loop unwindings considered; k-induction was performed on the instances marked with *)

B Additional Results for BMC Comparison

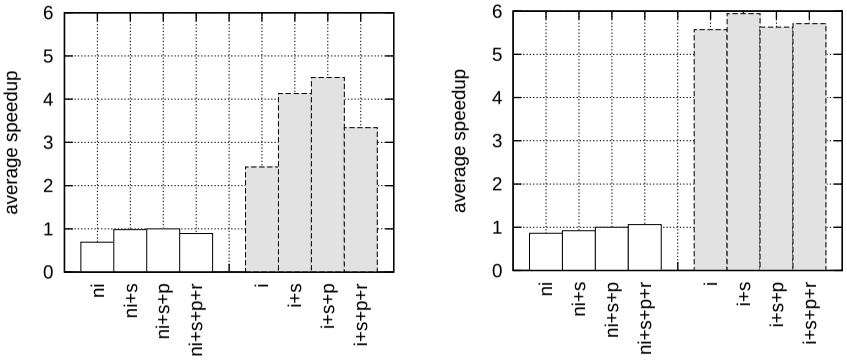

(a) SAT benchmarks (b) UNSAT benchmarks

Fig. 7: Effect of slicing, SAT formula preprocessing and bitvector refinement

C SystemC Benchmark Characteristics

		loops		
name	LOC	bounded	unbounded	unwindings
bist_cell_sat.cil.c	240	0	2	10
kundu_sat.cil.c	290	3	2	10
mem_slave_tlm.1_sat.cil.c	724	0	13	10
mem_slave_tlm.2_sat.cil.c	729	0	13	10
mem_slave_tlm.3_sat.cil.c	734	0	13	10
mem_slave_tlm.4_sat.cil.c	739	0	13	10
mem_slave_tlm.5_sat.cil.c	744	0	13	10
pc_sfifo_1_sat.cil.c	172	0	4	10
pc_sfifo_2_sat.cil.c	214	0	4	10
pc_sfifo_3_sat.cil.c	258	0	4	10
pipeline_sat.cil.c	400	0	3	10
token_ring.01_sat.cil.c	210	0	5	10
token_ring.02_sat.cil.c	270	0	6	10
token_ring.03_sat.cil.c	330	0	7	10
token_ring.04_sat.cil.c	390	0	8	10
token_ring.05_sat.cil.c	450	0	9	10
token_ring.06_sat.cil.c	510	0	10	10
token_ring.07_sat.cil.c	570	0	11	10
token_ring.08_sat.cil.c	630	0	12	10
token_ring.09_sat.cil.c	690	0	13	10
token_ring.10_sat.cil.c	750	0	14	10
token_ring.11_sat.cil.c	810	0	15	10
token_ring.12_sat.cil.c	870	0	16	10
token_ring.13_sat.cil.c	930	0	17	10
toy_sat.cil.c	315	0	6	10
kundu1_unsat.cil.c	233	2	2	3
kundu2_unsat.cil.c	285	3	2	2
pc_sfifo_1_unsat.cil.c	173	0	4	1
pc_sfifo_2_unsat.cil.c	215	0	4	1
pipeline_unsat.cil.c	400	0	3	5
token_ring.01_unsat.cil.c	217	0	5	3
token_ring.02_unsat.cil.c	277	0	6	3
token_ring.03_unsat.cil.c	337	0	7	3
token_ring.04_unsat.cil.c	397	0	8	3
token_ring.05_unsat.cil.c	457	0	9	3
token_ring.06_unsat.cil.c	517	0	10	3
token_ring.07_unsat.cil.c	577	0	11	3
token_ring.08_unsat.cil.c	637	0	12	3
token_ring.09_unsat.cil.c	697	0	13	3
token_ring.10_unsat.cil.c	757	0	14	3
token_ring.11_unsat.cil.c	817	0	15	3
token_ring.12_unsat.cil.c	877	0	16	3
token_ring.13_unsat.cil.c	937	0	17	3
token_ring.14_unsat.cil.c	875	0	16	3
token_ring.15_unsat.cil.c	935	0	17	3
toy1_unsat.cil.c	317	0	6	3
toy2_unsat.cil.c	314	0	6	3
transmitter.01_unsat.cil.c	197	0	6	2
transmitter.02_unsat.cil.c	256	0	7	2
transmitter.03_unsat.cil.c	315	0	8	2
transmitter.04_unsat.cil.c	374	0	9	2
transmitter.05_unsat.cil.c	433	0	10	2
transmitter.06_unsat.cil.c	492	0	11	2
transmitter.07_unsat.cil.c	551	0	12	2
transmitter.08_unsat.cil.c	610	0	13	2
transmitter.09_unsat.cil.c	669	0	14	2
transmitter.10_unsat.cil.c	728	0	15	2
transmitter.11_unsat.cil.c	787	0	16	2
transmitter.12_unsat.cil.c	846	0	17	2
transmitter.13_unsat.cil.c	905	0	18	2
transmitter.15_unsat.cil.c	905	0	18	1
transmitter.16_unsat.cil.c	961	0	19	1

Table 3: SystemC benchmark characteristics (name of the benchmark, lines of code, number of bounded loops, number of unbounded loops, and number of loop unwindings considered)